MANCHESTER 1824
The University of Manchester

Artificial Intelligence: a disruptive technology

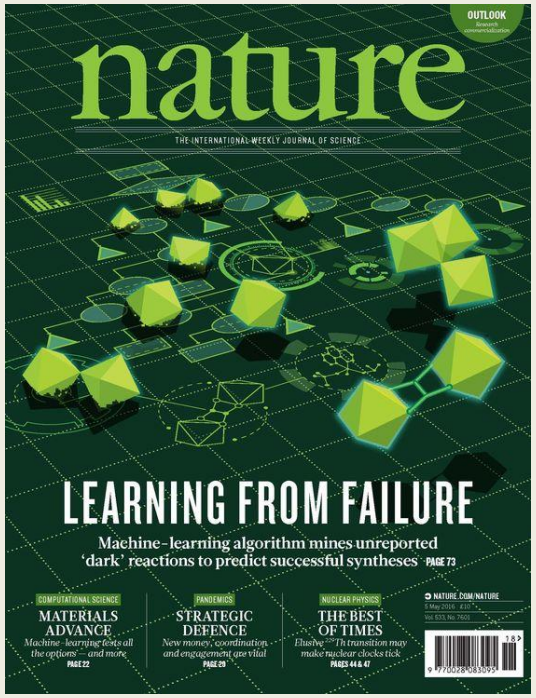

CAVE AT

MANCHESTER 1824
The University of Manchester

University, National & European initiatives:

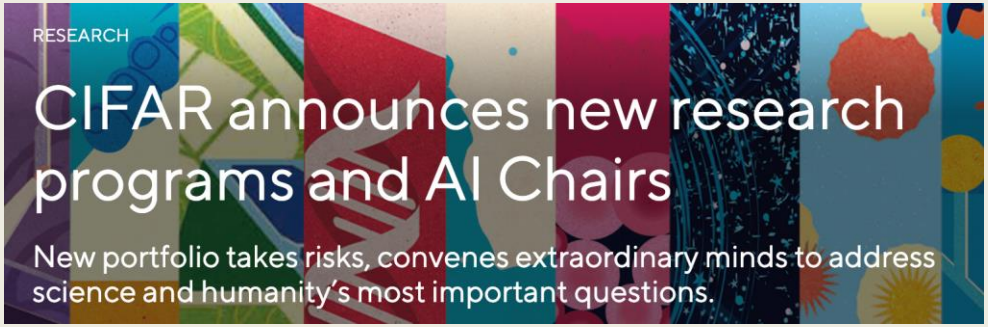

The Alan Turing Institute

The Alan Turing Institute

- Network of industry, charity, government partners
- Network of university members
- Strategic UK government investment

Goals:

- Innovate and develop world-class research
- Real-world problems
- Train the next generation
- Advising policy-makers and public

The Alan Turing Institute

University partners:

University of Cambridge

The University of Edinburgh

University of Oxford

UCL

Warwick
The University of Warwick

University of Bristol

University of Birmingham

University of Exeter

University of Leeds

Manchester 1824
The University of Manchester

Queen Mary University of London

Newcastle University

University of Southampton

The Alan Turing Institute

Strategic partners:

The Alan Turing Institute

Make algorithmic systems fair, ethical, and transparent

Design computers for the next generation of algorithms

Revolutionise healthcare

Deliver safer, smarter engineering

Supercharge research in science and humanities

Manage security in an insecure world

Shine a light on our economy

Foster government innovation

The Alan Turing Institute

Deliver safer, smarter engineering

Impact Story

Bridging the gap between physical and digital

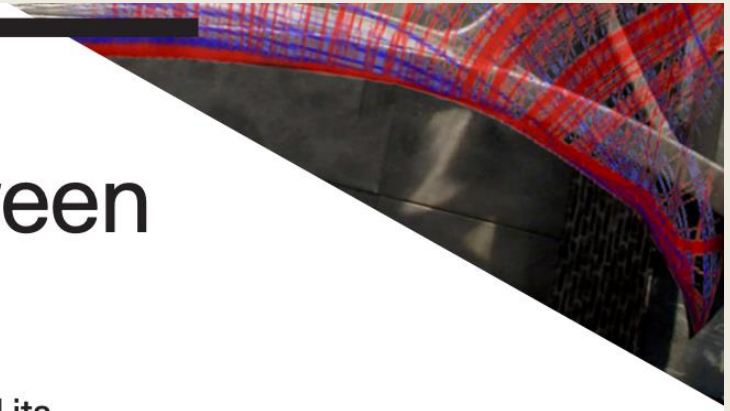

The Turing's data-centric engineering programme and its collaborators are unlocking insights into the world-first 3D-printed steel bridge, using innovative data science techniques and 'digital twin' technology.

- Dutch 3D-printing company MX3D in conjunction with a huge number of industrial and academic collaborators, unveiled the bridge at October 2018 Dutch Design Week.
- Bridge is a living laboratory for data scientists. Turing and collaborators have been developing a 'digital twin' of the bridge to help analyse data from sensors fixed to it.
- Also developing a long-term structural health monitoring network to ensure that the bridge remains safe during its lifetime.
- Testing has been supplemented by Turing researchers using statistical approaches to understand the "intrinsically random" properties of printed steel.
- Bridge will be installed over a canal

Impact

- Aided the development and analysis of the first 3D-printed steel bridge
- Fostered new collaborations between fields and researchers
- Allowing City of

The Alan Turing Institute

Revolutionise healthcare

SCIENTIFIC REPORTS

OPEN

Prognostication and Risk Factors for Cystic Fibrosis via Automated Machine Learning

Received: 24 January 2018
Accepted: 3 July 2018
Published online: 26 July 2018

Ahmed M. Alaa[1] & Mihaela van der Schaar[2,3,1]

Accurate prediction of survival for cystic fibrosis (CF) patients is instrumental in establishing the optimal timing for referring patients with terminal respiratory failure for lung transplantation (LT). Current practice considers referring patients for LT evaluation once the forced expiratory volume (FEV_1) drops below 30% of its predicted nominal value. While FEV_1 is indeed a strong predictor of CF-related mortality, we hypothesized that the survival behavior of CF patients exhibits a lot more heterogeneity. To this end, we developed an algorithmic framework, which we call AutoPrognosis, that leverages the power of machine learning to automate the process of constructing clinical prognostic models, and used it to build a prognostic model for CF using data from a contemporary cohort that involved 99% of the CF population in the UK. AutoPrognosis uses Bayesian optimization techniques to automate the process of configuring ensembles of machine learning pipelines, which involve imputation, feature processing, classification and calibration algorithms. Because it is automated, it can be used by clinical researchers to build prognostic models without the need for in-depth knowledge of machine learning. Our experiments revealed that the accuracy of the model learned by AutoPrognosis is superior to that of

The Alan Turing Institute

Make algorithmic systems fair, ethical, and transparent

Why a Right to Explanation of Automated Decision-Making Does Not Exist in the General Data Protection Regulation

Sandra Wachter*, Brent Mittelstadt** and Luciano Floridi***

Key Points

- Since approval of the European Union General Data Protection Regulation (GDPR) in 2016, it has been widely and repeatedly claimed that a 'right to explanation' of all decisions made by automated or artificially intelligent algorithmic systems will be legally mandated by the GDPR once it is in force, in 2018.
- However, there are several reasons to doubt both the legal existence and the feasibility of such a

as well as the significance and the envisaged consequences of automated decision-making systems, what we term a 'right to be informed'.

- The ambiguity and limited scope of the 'right not to be subject to automated decision-making' contained in Article 22 (from which the alleged 'right to explanation' stems) raises questions over the protection actually afforded to data subjects.
- These problems show that the GDPR lacks precise language as well as explicit and well-defined rights and safeguards against automated decision-making, and therefore runs the risk of being toothless.

The Alan Turing Institute

Make algorithmic systems fair, ethical, and transparent

This is very difficult in general – a societal rather than technical challenge

Tutorial: 21 fairness definitions and their politics

Arvind Narayanan

Update: this tutorial was presented at the Conference on Fairness, Accountability, and Transparency, Feb 23 2018. Watch it here.

Computer scientists and statisticians have devised numerous mathematical criteria to define what it means for a classifier or a model to be fair. The proliferation of these definitions represents an attempt to make technical sense of the complex, shifting social understanding of fairness. Thus, these definitions are laden with values and politics, and seemingly technical discussions about mathematical definitions in fact implicate weighty normative questions. A core component of these technical discussions has been the discovery of trade-offs between different (mathematical) notions of fairness; these trade-offs deserve attention beyond the technical community.

MANCHESTER 1824
The University of Manchester

The Alan Turing Institute

What can and should universities do as centres of research and as institutions educating the generations that will live in a world with artificial intelligence everywhere?

- Promote responsible research & innovation
- Provide ethical leadership by engaging with policymakers and other stakeholders
- Educate everyone about AI: benefits, risks & trade-offs
- Train developers and users
- Do all these things through national and international collaborations

MANCHESTER 1824
The University of Manchester

Educate everyone about AI: benefits, risks & trade-offs

University College for Interdisciplinary Learning

Units

2018/19 semester 2

AI

The Art of Enterprise

Bioethics

Crisis of Nature

Current Topics in Biology

Digital Society

Essential Enterprise

From Frankenstein to the Matrix

The Information Age

Introduction to British Sign Language - Part 1

Introduction to British Sign Language Part 2

Leadership in Action

AI: Robot Overlord, Replacement or Colleague?

Unit Code

UCIL20122

Unit Details

- Level 2
- 10 Credits
- Schools of Computer Science and Health Sciences

Overview

Artificial intelligence (AI), the ability of machines to learn from data, make decisions and perform actions, is now creeping into every aspect of our lives. This unit explores the mechanisms, implications and ethics of an environment where AI plays an increasingly important role.

- We will consider the science behind the headlines to help you develop an informed opinion regarding the complexities of the use of AI in society
- We will discuss the conceptual frameworks behind AI methodologies and the sources of the data on which they operate
- We will provide an introduction to computational thinking what sort of problems can AI realistically be expected to help with?
- There will be an in depth analysis of a series of case

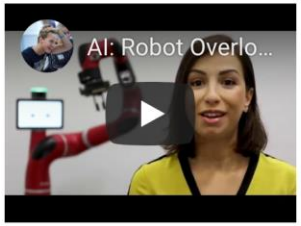

"I thought that online learning would be limited to a series of lecture slides in Blackboard. It was so much more than that. The content and learning experience was much better than I had expected! It was interactive and engaging with the added bonus that I could study at my own pace in bite-sized chunks. There was also lots of support from online tutors. I would recommend online learning to anyone."